

**In the Claims:**

Please amend claims 1, 2, 20, 21, 33 and 34, please cancel claims 12 and 32, and please add claims 37 and 38, as indicated below. A complete listing of claims follows.

1. (Currently amended) An apparatus, comprising:

a memory configured to provide an addressable block operand storage space and to store block operands within the block operand storage space;

a functional unit configured to perform a block operation on one or more of the block operands to generate a block result; and

a cache accumulator memory coupled to the memory and the functional unit, wherein the cache accumulator memory comprises a plurality of block storage locations, wherein the cache accumulator memory is configured to receive a set of one or more instructions to perform a first accumulation operation, wherein a first instruction in the set uses a first address in the memory to identify a first block operand;

~~wherein the cache accumulator memory is configured as a cache of the memory;~~

~~wherein the cache accumulator memory is configured~~ plurality of block storage locations are configured to cache a portion of the block operand storage space of the memory and to accumulate an intermediate result of the first accumulation operation, wherein the intermediate result is both a result of and an operand of the first accumulation operation, such that during the first accumulation operation, the plurality of block storage locations are concurrently configured both to cache certain ones of the block operands and to accumulate the intermediate result of the first accumulation operation; and

wherein in response to receiving the first instruction in the set, the cache accumulator memory is configured to access an associativity list comprising an indication that a first set of the block storage locations is allocated to the first accumulation operation and, in response to the indication, to provide the first block operand to the functional unit from the first set of block storage locations and to store the block result generated by the functional unit into the first set of block storage locations.

2. (Currently amended) The apparatus of claim 1, wherein the cache accumulator memory comprises a dual-ported memory configured to provide the plurality of block storage locations.

3. (Previously presented) The apparatus of claim 1, wherein the cache accumulator memory comprises at least two independently interfaced memory banks, wherein the cache accumulator memory is configured to provide the first block operand from a first block storage location in a first one of the independently interfaced memory banks and to store the block result in a second block storage location in a second one of the independently interfaced memory banks, wherein the first set of block storage locations comprises the first block storage location and the second block storage location.

4. (Previously presented) The apparatus of claim 1, wherein the cache accumulator memory is configured to indicate whether a particular block operand stored in the cache accumulator memory is modified with respect to a copy of that particular block operand in the memory.

5. (Previously presented) The apparatus of claim 1, wherein the cache accumulator memory is configured to load a copy of the first block operand into the first set of block storage locations in the cache accumulator memory from the memory in

response to the first block operand not being present in the cache accumulator memory when the first instruction is received.

6. (Previously presented) The apparatus of claim 5, wherein if all of the block storage locations in the cache accumulator memory are currently storing valid data when the first instruction is received, the cache accumulator memory is configured to select the first set of block storage locations and to load the copy of the first block operand into the first set of block storage locations, wherein the cache accumulator memory is further configured to update the indication in the associativity list to indicate that the first set of the block storage locations is allocated to the first accumulation operation in response to selecting the first set of block storage locations.

7. (Previously presented) The apparatus of claim 6, wherein the cache accumulator memory is configured to use a least recently used algorithm to select the first set of block storage locations to overwrite.

8. (Previously presented) The apparatus of claim 6, wherein if data to be overwritten in the first set of block storage locations is modified with respect to a copy of that data in the memory, the cache accumulator memory is configured to write the data back to the memory before loading the copy of the first block operand into the first set of block storage locations.

9. (Previously presented) The apparatus of claim 5, wherein in response to loading the first block operand into the first set of block storage locations, the cache accumulator memory is configured to update a tag associated with the first set of block storage locations to indicate that the first block operand is stored within.

10. (Previously presented) The apparatus of claim 1, wherein the cache accumulator memory is configured to update the associativity list in response to storing the block result generated by the functional unit, wherein the cache accumulator memory is configured to update the associativity list by updating a tag associated with the first set

of block storage locations to indicate that the block result is stored within the first set of block storage locations.

11. (Previously presented) The apparatus of claim 10, wherein the cache accumulator memory is configured to update the tag by setting the tag to equal a first portion of address bits of a second address in the memory that identifies the block result.

12. (Canceled)

13. (Original) The apparatus of claim 1, wherein the functional unit is configured to perform a parity calculation on the block operand.

14. (Original) The apparatus of claim 1, wherein the operation comprises a parity calculation, and wherein the command is issued by a storage system controller.

15. (Original) The apparatus of claim 14, wherein the functional unit is configured to calculate a parity block from a plurality of data blocks in a stripe of data, wherein the first block operand is a first one of the data blocks in the stripe of data.

16. (Original) The apparatus of claim 1, wherein the functional unit is configured to perform the operation on two block-operands.

17. (Previously presented) The apparatus of claim 16, wherein a first of the two block-operands is the first block operand stored in the cache accumulator memory and a second of the two block-operands is provided on a data bus coupled to provide operands to the functional unit.

18. (Previously presented) The apparatus of claim 16, wherein a first of the two block-operands is the first block operand stored in the cache accumulator memory and a second of the two block-operands is provided from the memory.

19. (Previously presented) The apparatus of claim 1, wherein the cache accumulator memory is configured to store a word of the block result during an access cycle in which cache accumulator memory also provides a word of the first block operand to the means for performing a block operation.

20. (Currently amended) A method of performing a block accumulation operation using a cache accumulator memory that comprises a plurality of block storage locations, the method comprising:

receiving a first command in a set of commands used to implement an accumulation operation, wherein the first command is an instruction to perform an operation on a first block operand identified by a first address in a memory and to store a result of the operation, wherein the result is identified by a second address in the memory, wherein the memory is configured to provide an addressable block operand storage space and to store block operands including the first block operand within the block operand storage space, and wherein the plurality of block storage locations are configured to cache a portion of the block operand storage space of the memory ~~the cache accumulator memory is configured as a cache of the memory;~~

in response to said receiving a first command:

accessing an associativity list comprising an indication that a first set of block storage locations of the cache accumulator memory is allocated to the first accumulation operation;

in response to the indication, providing the first block operand from the first set of block storage locations to a functional unit and storing a block result of the operation generated by the functional unit into the first set of block storage locations;

wherein the ~~cache accumulator memory~~ is plurality of block storage locations are configured to accumulate an intermediate result of the first accumulation operation, wherein the intermediate result is both a result of and an operand of the first accumulation operation, such that during the accumulation operation, the plurality of block storage locations are concurrently configured both to cache certain ones of the block operands and to accumulate the intermediate result of the accumulation operation.

21. (Currently amended) The method of claim 20, wherein the cache accumulator memory comprises a dual-ported memory configured to provide the plurality of block storage locations, wherein said storing comprises overwriting the first block operand with the block result.

22. (Previously presented) The method of claim 20, wherein the cache accumulator memory comprises at least two independently interfaced memory banks, wherein said loading comprises loading the first block operand into a first block storage location in a first one of the independently interfaced memory banks and wherein said storing comprises storing the block result in a second block storage location in a second one of the independently interfaced memory banks, wherein the first set of block storage locations comprises the first block storage location and the second block storage location.

23. (Original) The method of claim 20, further comprising selecting the first set of block storage locations and loading the first block operand into the first set of block storage locations if all of the block storage locations are currently storing valid data when the first command is received.

24. (Original) The method of claim 23, wherein said selecting comprises using a least recently used algorithm to select the first set of block storage locations.

25. (Original) The method of claim 23, further comprising writing data in the first set of block storage locations back to the memory if the data is modified with respect to a copy of that data in the memory.

26. (Original) The method of claim 20, further comprising updating the indication in the associativity list by updating a tag associated with the first set of block storage locations to indicate that the block result is stored within the first set of block storage locations in response to storing the block result generated by the functional unit.

27. (Original) The method of claim 26, wherein said updating the tag comprises setting the tag to equal a first portion of address bits of the second address.

28. (Original) The method of claim 20, further comprising the functional unit performing a parity calculation on the first block operand to generate the block result in response to said providing.

29. (Original) The method of claim 20, wherein the operation comprises a parity calculation, and wherein the command is issued by a storage system controller.

30. (Original) The method of claim 20, further comprising the functional unit performing the operation on the first block operand and a second block operand in response to said providing.

31. (Original) The method of claim 30, further comprising a data bus providing the second of the two block operands to the functional unit.

32. (Canceled)

33. (Currently amended) An apparatus, comprising:

means for storing data configured to provide an addressable block operand storage space and to store block operands within said block operand storage space;

means for performing a block operation on one or more of the block operands to generate a block result; and

means for storing the block result, wherein the means for storing the block result are coupled to the means for storing data and the means for performing a block operation, wherein the means for storing the block result comprise a plurality of block storage locations, wherein the means for storing the block result receive a first instruction comprised in a set of one or more instructions to perform a first accumulation operation, wherein the first instruction uses a first address in the means for storing data to identify a first block operand;

wherein the ~~means for storing the block result is~~ plurality of block storage locations are configured to cache a portion of the block operand storage space ~~configured as a cache of the means for storing data~~;

wherein the ~~means for storing the block result is~~ plurality of block storage locations are configured to accumulate an intermediate result of the first accumulation operation, wherein the intermediate result is both a result of and an operand of the block accumulation operation, such that during the first accumulation operation, the plurality of block storage locations are concurrently configured both to cache certain ones of the block operands and to accumulate the intermediate result of the first accumulation operation; and

wherein in response to the first instruction, the means for storing the block result access an associativity list that comprises an indication that a first set of



the block storage locations is allocated to the first accumulation operation, wherein in response to the indication, the means for storing the block result provide the first block operand from the first set of block storage locations to the means for performing the block operation and store the block result in the first set of block storage locations.

34. (Currently amended) A data processing system, comprising:

a host computer system;

a storage array;

an interconnect coupled to the host computer system and the storage array and configured to transfer data between the host computer system and the storage array; and

a parity calculation system configured to perform parity operations on data stored to the storage array, wherein the parity calculation system comprises a memory configured to provide an addressable block operand storage space and to store block operands within the block operand storage space, a cache accumulator memory, and a parity calculation unit;

wherein the cache accumulator memory comprises a plurality of block storage locations and is configured to receive a set of one or more instructions to perform a first accumulation operation, wherein a first instruction in the set uses a first address in the memory to identify a first one of the block operands;

~~wherein the cache accumulator memory is configured as a cache of the memory;~~

wherein the ~~cache accumulator memory is configured~~ plurality of block storage locations are configured to cache a portion of the block operand storage space of the memory and to accumulate an intermediate result of the first accumulation operation, wherein the intermediate result is both a result of and an operand of the first accumulation operation, such that during the first accumulation operation, the plurality of block storage locations are concurrently configured both to cache certain ones of the block operands and to accumulate the intermediate result of the first accumulation operation; and

wherein in response to receiving the first instruction in the set, the cache accumulator memory is configured to access an associativity list comprising an indication that a first set of the block storage locations is allocated to the first accumulation operation and, in response to the indication, to provide the first block operand to the parity calculation unit from the first set of the block storage locations and to store the block result generated by the parity calculation unit into the first set of block storage locations.

35. (Previously presented) The data processing system of claim 34, wherein the parity calculation unit is configured to perform a parity calculation on the first block operand provided by the cache accumulator memory and a second block operand provided on a data bus.

36. (Original) The data processing system of claim 31, wherein the parity calculation system is configured to calculate a parity block from a plurality of data blocks in a stripe of data when performing the first accumulation operation, wherein the first block operand is a first one of the data blocks in the stripe of data and wherein the second block operand is a second one of the data blocks in the stripe of data.

37. (New) The apparatus as recited in claim 1, wherein the block operands stored by the memory within the block operand storage space are received from one or more mass storage devices.

38. (New) The method as recited in claim 20, wherein the block operands stored by the memory within the block operand storage space are received from one or more mass storage devices.